

Beyond intelligent vacuum cleaners

AD-A280 565



JUN 24 1994

S F D

Don Brutzman
Naval Postgraduate School
Code OR/Br

Monterey California 93943-5000 USA
brutzman@nps.navy.mil
(408) 656-2149 work, (408) 656-2595 fax

94-19167



32

Abstract

Designing the ultimate autonomous vacuum cleaner is the quantifiable and clearly-defined goal of this symposium. It falls short of addressing the problems which most interest roboticists working in artificial intelligence. The painstaking progress and restricted performance of current mobile robots apparently necessitates such a low-level review. Overcoming such challenges and setbacks is essential if robots are to become capable of operating in more interesting regimes. Several questions are presented to help evaluate whether a robotics research program remains oriented towards meaningful long-term goals of autonomy, intelligence and capable performance.

1 Motivation

Mobile robotics is a field of exciting challenge particularly for artificial intelligence (AI) researchers. The practice of mobile robotics is also fraught with frustration. Despite years of tremendous effort, progress has been slow. Working robots exhibit only limited autonomy and rudimentary skills. Current robot demonstrations do not approach the promise of AI theory. Few roboticists are satisfied with the state of the art.

Instantiating autonomous real-world agents is crucial for the fields of mobile robotics and AI. Given that current robots fall short of expectations and the bottleneck causes for that lack of rapid progress remain unclear, deliberately focusing on exhaustive and definitive solutions to the autonomous vacuum cleaner problem unfortunately appears to be a valid exercise in reassessment and consensus building. Nevertheless our long-standing preoccupation with current problems ought not to divert our attention from where we want to go. This essay attempts to ask questions that may help lift our focus past today's problems towards the horizon of our long-range goals.

Published in the American Association for Artificial Intelligence (AAAI) Fall Symposium on Applications of Artificial Intelligence for Instantiating Real-World Agents, Raleigh, North Carolina, October 22-24, 1993, pp. 23-25.

This document has been approved
for public release and sale; its
distribution is unlimited.

2 Working on the wrong problems

The most sophisticated vacuum cleaner in the world picks up trash. Your mother can quickly tell you that picking up trash is a well-understood problem. As researchers we ought to be very unhappy that we still need to spend a few days wrestling with the pros and cons of various trash collection methodologies. Hopefully workshop participants will be able to conclude that this problem is solvable theoretically, practically and economically. We need to declare victory and demonstrate success on numerous lower-level robotics problems in order to proceed to the more interesting challenges of real-world autonomy.

Dissatisfaction with current robotics progress is a valuable attitude because years of preoccupation with specific lower-level problems can lead to an inflated opinion of their significance. Certainly many of the component problems associated with an autonomous vacuum cleaner are difficult. So what? Get a working solution and move up to higher problems. Further refinements are primarily engineering. Perennially working on the same low-level problems without satisfactory resolution may be an indicator that your efforts are being directed against the wrong problem.

What's a wrong problem? A wrong problem is one that keeps us away from our avowed goals of understanding and instantiating intelligence. The absence of anything like the android _____ (your favorite fictional robot here) should be considered a standing indictment that AI and mobile robotics are very far from keeping pace with our imaginations. Our fields appear unable to come anywhere near replicating an entity with physical or mental attributes comparable in scope to human capabilities. Lots of work remains! We can't afford to squander efforts by repeatedly reinventing narrow solutions to narrow problems.

"If you do not work on an important problem, it's unlikely that you'll do important work" (Hamming 1986). Working on wrong problems is a cul-de-sac since the best you can hope for (their solution) is by definition of little importance. Working on right problems keeps your goals pointed at the horizons of AI. Working on important problems also provides a greatly clarifying context which puts everything into perspective.

94 6 23 021

DTIC QUALITY INSPECTED 2

Here is an example of how your choice of problem can add clarity. While at this symposium you become exposed to a number of new and fascinating ideas. Choose your preferred reaction:

- a) "How will this idea help my robot vacuum trash?"
- b) "How will this idea help my robot accomplish something that's never been done before?"

3 Working on the right problems

Choice of problem has profound consequences. Good examples of this point abound in a critical branch of robotics: interpretation of sensor data. Too little data enables only limited responses. Too much data can quickly overload interpretation and decision-making processes. In most robots, choice of sensor and processing architecture directly define and constrain the robot's capabilities (Miller 1993). The choice of sensor problem to pursue may bound the scope of possible robot interactions, just as choice of vocabulary can bound the scope of expressible thoughts.

Typical problem choices often lead to a robot that is anthropomorphic in task domain. Building an automatic vacuum cleaner is one such exercise. Navigating an office environment is another. It seems that most robot problem domains are identical to problems solved by people. This is not surprising since easy access to robots enables us to better test and develop them. However anthropomorphism of task can be an unnecessary restriction. It may on occasion be more revealing to focus on unfamiliar problems that have no direct human corollary, such as exploration and interaction in remote hazardous environments.

Similarly, anthropomorphism of robot "thought" may be an unnecessary restriction when choosing the right problem domain for a robot. While many AI paradigms are modeled after human problem-solving techniques, numerous other AI paradigms are not. Machine learning methodologies are particularly prone to achieve solutions quite dissimilar from those a person might develop. The intelligent components we incrementally enable in our robots are likely to form a new vocabulary which will lead to still different thought patterns and emergent behaviors. We should not intentionally restrict robotics research to only those problems which seem humanly solvable.

"There are sounds that you cannot hear but that some animals, such as dogs, can hear. There are tastes similarly that you cannot taste, and sights you cannot see, being limited to about one octave of the electromagnetic spectrum. These limitations come from the sense organs you have and not from the source. Since this is true then why not unthinkable thoughts - thoughts that you cannot think, given the way your mind is wired? Why should we blandly assume that we can think anything?" (Hamming 1990)

4 What are the right questions?

Here are some questions you might ask periodically to help your robot develop autonomous capabilities.

What is the robot doing right now? Let it run all the time. A robot has to be safe, reliable and self-sufficient to run continuously - this is autonomous self-preservation. As it operates it will encounter new situations which result in success or failure. Every time the robot breaks or stops unexpectedly you learn something. Every fix hopefully improves robot autonomy. A significant additional benefit is that the longer it runs, the more opportunities it has to perform machine learning.

What is the robot "seeing" right now? If we can "look" through the robot's sensors then we can perceive the world from its perspective. This will allow us to better understand the implications of the arbitrary sensor constraints we have imposed. If the sensor is dissimilar to human senses (e.g. laser, sonar, radar, odometry) then it is imperative that we develop an intuitive feel for that sensor's capabilities and weaknesses. A further benefit is that we become expert in employing autonomous sensors, which is particularly useful for remote monitoring of hazardous, expensive or inaccessible locations.

Run that by again? We will undoubtedly observe some unexpected results. Repeatability (or at least playback) of robot behavior is essential to clarify our understanding of those results. A replay capability not only permits testing of theory but also validation and verification of our models. Repeatability and playback are also valuable for educating others and broadening access to our robots.

How are we using the sensors? Keep in mind that sensor inputs define the world view. Have you thought through the relationships between your sensors, your knowledge representations, and the ultimate problems your robot is intended to solve? Here is a sea story about sensors. During a previous tour as a submarine officer I had the opportunity to ride a

number of different submarines during operational testing of a new torpedo design. As a rider I periodically had to brief the commanding officer (CO) regarding exercise status. Some COs were notoriously difficult to locate. This was unusual since there are only so many places you can go on a submarine. After some time I finally noticed something about those guys: (1) they were top-notch tacticians; (2) when no one knew where to find them, they were in the sonar shack. The COs with the best world view knew the capabilities and limitations of their sensors inside out.

5 What are the right metaquestions?

Here are some questions you might ask periodically to help ensure your robotic efforts are working on the right problems.

What are the important problems? This question seems fundamental, and it has a critical follow-on question:

Why aren't I working on them? If your answers are not satisfying it is time to consider changing course.

Does this component solution fit our paradigm? This is a dangerous question. Be wary of architectural dogmas that constrain your options. There are many different things in the world that a robot needs to deal with. Plan and build for diverse approaches. Hierarchical abstraction and hybrid approaches will be needed to avoid pathological deadlocks while accommodating diversity. As an example, we don't look at another person and think "my what a highly tuned collection of approximately 10,000 cooperative behaviors!" Rather we say "Hi Terri!"

What is the real problem? Seemingly-intractable problems sometimes just need to be looked at from a completely different perspective to be solved. AI folks are usually pretty good at this.

Is it grounded? Not taking any theory or implementation as gospel until you've run it and tried to break it is a pretty good rule of thumb. Robotics is empirical.

Where's the AI? (Shank 1992) If you can't answer this question, then what is it really that your robot is doing? Shank's essay deserves repeated reading. Robotics discussions (this essay included) often sidestep the philosophically-controversial term "intelligence" by using the more observable characteristic of autonomy. No matter. Intelligent behavior remains the ultimate goal. In mobile robotics you don't have to work very hard to prove your robot's intelligence because it ought to be obvious. If our robots don't appear "smart" to our kids and our sponsors then we lose interest, lose funding and find other employment.

6 Recommendations

Researchers in mobile robotics ought to be dissatisfied with the less-than-breathtaking performance of robots today. A questioning attitude is essential to keep robot efforts focused on meaningful progress instead of isolated piecemeal successes.

When you look at stacks of journals and books in your office or library, you are not looking at many problems but rather many solutions. Thousands of individual component solutions to mobile robotics problems exist today. However very few are actually implemented in working robots. Preoccupations with long-standing low-level problems need to yield to broader higher-level concerns. Basic solutions need to be plugged together to provide robust working robots as AI research platforms. Mobile robotics needs to move up to the real challenges of autonomy, intelligence, and practical instantiation of capable real-world agents.

Acknowledgments

I thank Richard W. Hamming for numerous ideas that significantly contributed to this essay. Thanks also to David P. Miller and Michael J. Zyda for suggestions and support.

References

- Hamming, Richard W. 1986. "You and Your Research," Bell Communications Research Colloquium Seminar, Morristown, New Jersey, delivered 7 March 1986, p. 8.
- Hamming, Richard W. 1990. "The Future of Science and Engineering," Superintendent's Guest Lecture, Naval Postgraduate School, Monterey, California, delivered 19 April 1990, p. 10.
- Miller, David P. 1993. "Cleaning Up With Second-Hand Sensing," AAAI Fall Symposium on Instantiating Real-World Agents, Raleigh, North Carolina, October 22-24, 1993, pp. 99-100.
- Shank, Roger C. 1991. "Where's the AI?" AI Magazine, vol. 12 no. 4, Winter 1991, pp. 38-49.

Accession For	
NTIS CRAN	
DTIC IAF	
Union of	
Justification	
By	
Distribution	
Availability	
Dist	
A-1	